Quantifying the Diaspora of Knowledge in the Last Century

Manlio De Domenico[1], Elisa Omodei[2], Alex Arenas[1]
[1]Departament d'Enginyeria Informàtica i Matemàtiques, Universitat Rovira i Virgili, 43007 Tarragona, Spain

Academic research is driven by several factors causing different disciplines to act as "sources" or "sinks" of knowledge. However, how the flow of authors' research interests – a proxy of human knowledge – evolved across time is still poorly understood. Here, we build a comprehensive map of such flows across one century, revealing fundamental periods in the raise of interest in areas of human knowledge. We identify and quantify the most attractive topics over time, when a relatively significant number of researchers moved from their original area to another one, causing what we call a "diaspora of the knowledge" towards sinks of scientific interest, and we relate these points to crucial historical and political events. Noticeably, only a few areas – like Medicine, Physics or Chemistry – mainly act as sources of the diaspora, whereas areas like Material Science, Chemical Engineering, Neuroscience, Immunology and Microbiology or Environmental Science behave like sinks.

arXiv:1604.00696v1 [cs.SI] 3 Apr 2016

I. INTRODUCTION

Nowadays, the research carried out by academics in all areas of human knowledge is heavily driven by exogenous factors, such as allocation of funding resources or political interests [1, 2]. Two decades ago, pioneering studies by Etzkowitz and Leydesdorff already put in evidence the importance of relationships between university, industry and government [3–5], a "triple helix" that shapes and drives the development of knowledge, impelling researchers to change research interests or their institution [6–8]. The structure and evolution of human knowledge has been extensively investigated by observing, for instance, how academics tend to choose their co-authors or they physically move between different research institutions [3–12]. These analyses, often based on citation patterns among authors, institutions, papers or journals, allow to understand how disciplines are related to each other in terms of scientific production and impact, but are not intended to quantify the flow of knowledge in science or to identifying crucial periods for the development of human knowledge. In fact, the interest of researchers are not rarely driven by currently available funding opportunities or by political choices, an emblematic example being the investments in nuclear physics during the World War II. Such factors, often external to the context of academy research, act as catalysts pushing researchers to leave their current area of interest towards different areas.

To study this phenomenon of "knowledge diaspora", we consider the Microsoft Academic Graph, a data set of more than 35,000,000 of papers published in more than 21,000 different journals in the last 100 years. We are interested in exploiting metadata information to classify each paper into one or more disciplines. Unfortunately, our exploratory analysis of the classification scheme released with the dataset, based on paper keywords, revealed some relevant drawbacks that would dramatically bias the more sophisticated analysis presented in this work. To cope with such limitations, we classified the papers according to the journal where they have been published, the rationale behind this choice being that journals tend to publish research studies that are, in general, more pertinent to their specific topic(s). For instance, it is difficult to publish in a physics journal a paper about humanities or biology, if this paper does not provide some physical insights that would make it suitable for an audience of physicists. Therefore, each journal is classified into one or more topics, fine-grained representations of academic knowledge, and into one or more areas, coarse-grained representations of academic knowledge. We use the SCImago classification, where there are 306 unique topics grouped into 27 distinct areas of knowledge to assign topics and areas to each paper, according to its journal. One possible cause of criticism might be that such classification is too recent to characterize adequately journals existing at the beginning of the past century. However, it must be remarked that we are focusing our attention on those journals that have a long-established tradition – i.e., from a few decades up to one century – and are unlikely to have dramatically changed their area of reference across years. We have divided the data set into non-overlapping temporal snapshots of 5 years, from 1910 to 2014. A snapshot marked with a year refers to a period between that year and 4 years later, e.g. 2000 refers to the period 2000-2004. To trace the changes in research interests of every author in the data set from one temporal snapshot to the successive, we count how many authors published in topic A at time τ and in the same or a different topic B at time $\tau + \Delta\tau$ (see Appendix). The volume of authors linking topics defines an evolving network of connections among topics, i.e. a multilayer (time-varying, weighted and directed) network [13–16]. The same procedure has been also applied to the coarser level of areas. The structure of these dynamical multilayer networks encode the publishing temporal dynamics of academics who change their research interests across knowledge topics and areas, respectively. In the following we will simply refer to these structures using the term network, avoiding to specify that they are time-varying and multilayer.

II. OVERVIEW OF THE DATA SET

The Microsoft Academic Graph is a heterogeneous graph containing scientific publication records, citation relationships between those publications, as well as authors, institutions, journals and conference "venues" and fields of study [17]. We used the latest publicly available updated version (31 August 2015) of this data set[1] in our study. However, our careful inspection of the data did not allow us to use the accompanying classification of papers into fields of study. The first obstacle was the number of different keywords classifying the papers: tens of thousands of categories providing a scheme too fine-grained for our study. A reduction of such keywords into more general topics would require machine learning and heuristics that would introduce other uncontrollable bias in the resulting classification. The second obstacle was the unclear mechanisms adopted to assign one or more keywords to each paper. In fact, we have found many misclassified papers, an emblematic case being a paper about Agricultural Science that has been classified in several topics, among which General Relativity. Instead, we gathered data from an external (publicly available) source. More specifically, we used SCImago Journal and Country Rank in 2014[2] to classify journals into 306 distinct research topics and 27 unique knowledge areas. Successively, we filtered out from the Microsoft Academic Graph data set all the papers that were not published in journals, thus excluding other venues such as conferences, and in particular we filtered out those papers published in journals that were not found in the SCImago classification. More than 35 millions of papers survived this filtering procedure, representing a promising 28.7% of the original data set, and more than 60% of the original set of papers published in journals only. The number of different journals matching the SCImago data set was 21,729, and we report in Supplementary Table I some information about the distribution of their multiplexity, i.e. the number of different topics and areas where they are classified. If a journal is classified into just one topic or area, the papers published in that journal will not be classified as multiplex, whereas, conversely, papers published in journals that are classified into more than one topic or area, will be treated as multiplex. Finally, it is worth remarking that we further reduced the dataset to avoid the effects of non-disambiguated authors. More specifically, we built the distribution of the number of papers per year of each author and we focused on the 99.9%-quantile distribution, i.e. we excluded the 0.1% of authors. This choice excluded all the names who authored more than 17 journal papers per year, the rational being that names with a higher number of papers per year probably corresponds to different authors having the same name.

[1] http://research.microsoft.com/en-us/projects/mag/
[2] http://www.scimagojr.com/journalrank.php

# of Topics	% of Journals
1	31.6%
2	33.4%
3	19.8%
4	9.2%
5	3.3%

# of Areas	% of Journals
1	50.9%
2	36.2%
3	9.8%
4	2.1%
5	0.5%

Table I. Multiplexity of journals with respect to topics and areas. We report the percentage of journals that are classified by SCImago in exactly 1, 2, ..., etc, topics (areas). Only statistics for the top five are reported, with rapidly decreasing percentage of journals classified in more than five topics (areas).

III. MULTILAYER NETWORK MODEL

The data set used in our study contains a huge amount of information about published papers and their authors. We focused on specific subsets of the data, including author name, the papers he/she published, the journal where they have been published and the publishing year. Thanks to the SCImago classification of areas of knowledge, we were able to assign one or more topics to each journal. Thus, we built a tripartite time-varying multilayer network $\mathcal{G}$ where for each temporal snapshot τ, a tripartite multiplex $\mathcal{M}$ is considered. Each multiplex is composed by layers $\mathcal{L}$ – identifying topics or areas of knowledge, depending on the application of interest – where there are three types of nodes: authors (A), papers (P) and journals (J). One or more authors are linked to the paper(s) they co-authored that, in turns, are linked to the journal where they have been published, resulting in a bipartite network linking nodes of type A to nodes of type P, and a bipartite network linking, at the same time, nodes of type P to nodes of type J. If a journal is classified in more than one topic or area, the links are replicated accordingly across layers. The resulting network is tripartite, because three types of nodes are involved, and multiplex, because nodes are replicated on different layers. For our purposes, we aggregated the tripartite network in each layer $l \in \mathcal{L}$ with respect to papers, in order to obtain multiplex bipartite networks of authors and journals only, for each temporal snapshot. Finally, each node is inter-connected to its replicas in other layers and temporal snapshots. The mathematical representation [14, 15] of $\mathcal{G}$ is a rank-6 tensor $G^{\alpha\tilde{\gamma}\bar{\epsilon}}_{\beta\tilde{\delta}\bar{\phi}}$, where indices $(\bar{\epsilon}, \bar{\phi})$ identify the temporal snapshot, $(\tilde{\gamma}, \tilde{\delta})$ identify the layers and (α, β) identify the nodes.

This complex network, however, is not the final object we worked with. In fact, our analysis is more focused on

changes in publication patterns across years. Mathematically, this means that we are more interested in the links between authors and journals exhibit between one temporal snapshot and the successive, i.e. in inter-layer links with respect to time. We derived a more suitable time-varying multilayer network $\mathcal{H}$ from $\mathcal{G}$ as follows. Let A_i be the $i-$th node of type A (i.e. authors) and J_k be the $k-$th node of type J (i.e. journals), regardless of topics (areas) classification and time. In $\mathcal{G}$, a link between A_i and J_k in layer l at time τ exists if $G^{il\tau}_{kl\tau} > 0$. Similarly, if in the successive snapshot $\tau' > \tau$ the same author A_i is linked to journal $J_{k'}$ (k' can be the same as k) in layer l' (l' can be the same as l), then $G^{il'\tau'}_{k'l'\tau'} > 0$. Clearly, an author might publish papers on different topics or areas at time τ but he/she will be, in general, more active on one or a few more. For this reason for each snapshot, we will consider only the layer where the author has been more active, i.e. where $G^{il\tau}_{kl\tau}$ is maximum with respect to l (note that if there is more than one layer where the author is equally active, we will consider all of those layers). We will indicate by $l^\star$ such layers. The components of the tensor representing $\mathcal{H}$ that encode inter-snapshot connections, are defined by

$$H^{il^\star\tau}_{il'^\star\tau'} = \Theta(G^{il^\star\tau}_{kl'^\star\tau}) \times \Theta(G^{il'^\star\tau'}_{k'l'^\star\tau'}), \tag{1}$$

i.e. an interconnection between an author at time τ and his/her replica at time $\tau' > \tau$ is present if and only if the author published at time τ and at time τ'. It is worth remarking that the replicas being linked are defined on layers $l^\star$ at time τ and $l'^\star$ at time τ', thus also connecting (possibly different) topics or areas across time. The presence of Heaviside step function $\Theta(\cdot)$ is to guarantee that each author is counted just once at this step, regardless if he/she produced more papers. It is evident that information about the flow of authors moving from one knowledge topic (or area) to another across time is only encoded in inter-snapshot connections among author's replicas, whereas the presence of journals as nodes is no more required, as well as intra-snapshot links, i.e. connections within the same temporal snapshot. Therefore, the tensor H representing $\mathcal{H}$ is defined on a smaller tensorial space with respect to G, because nodes are just authors instead of authors and journals. Moreover, it is also extremely sparse and, in fact, it can be further aggregated without loss of information, because of the absence of intra-snapshot links, by projecting the tensor into the space of topics (or areas) and time, getting rid of information about authors (see Appendix for details about this step). The resulting tensor $M^{\tilde{\gamma}\bar{\epsilon}}_{\tilde{\delta}\bar{\phi}}$, that is the one we used in our analysis, represents a multilayer network where nodes are topics (or areas), identified by indices $(\tilde{\gamma}, \tilde{\delta})$, and layers are temporal snapshots, identified by indices $(\bar{\epsilon}, \bar{\phi})$. Intra-layer links, i.e. connections among topics within the same temporal snapshot, are not present, whereas inter-layer links among topics encode the underlying flow of authors during consecutive periods of time.

IV. RESULTS

To gain the first insights about the knowledge diaspora across topics, we developed an ad hoc visualization (see Appendix) to put in evidence, for each topic, the intricate web of flows of authors incoming from and outgoing to other topics. We see in Fig. 1 a few emblematic cases corresponding to the diaspora observed in 1910, 1960 and in 2010, covering one century of academic publishing in all areas of knowledge. It is evident that one century ago authors were not contributing significantly outside their own area of expertise. After 50 years the diaspora is more prominent, with intense flows between topics of different areas, such as –Medicine– and –Biochemistry, Genetics and Molecular Biology–, between –Physics and Astronomy– and –Earth and Planetary Science–, or between –Chemistry– and –Chemical Engineering–. After 100 years, the diaspora is extremely evident, affecting basically all areas of knowledge.

The map of knowledge diaspora shown in Fig. 1 allows to get qualitative insight about this phenomenon, although it does not allow to quantify, for instance, the raise of research interest in specific topics. We will focus first our study on the emergence of topics of interest, by analyzing the variation of their incoming flows. To this aim, we quantify the attractiveness of a topic t through time $\delta_t(\tau)$, by tracking the evolution of the relative changes in the volume of authors $V_{tt'}(\tau)$ incoming from all other topics $t' \neq t$, at each temporal snapshot τ:

$$\delta_t(\tau) = \frac{1}{N_t - 1} \sum_{t' \neq t} \frac{V_{tt'}(\tau) - V_{tt'}(\tau - 5 \text{ years})}{V_{tt'}(\tau - 5 \text{ years})}, \tag{2}$$

being $N_t = 306$ the total number of topics considered. For each topic, it quantifies the average net relative change in the incoming flow. This parameter is sensitive to changes in the flow from one topic to another, even when this flow is rather small compared to the total incoming flow. Indeed, it might happen that a topic attracts a small flow of authors from many other topics or a huge flow of authors from a rather small set of other topics. The parameter $\delta_t(\tau)$ would detect both patterns and assign a similar score in the two cases. Other aggregated parameters, such as the relative change in the overall incoming flow per topic, are not able to capture this type of patterns, that would be inevitably hidden by larger flows with possibly less significant relative variations over time.

For each snapshot τ separately, we look for the most attractive topic, the one with the highest value of $\delta_t(\tau)$. The results, shown in Fig. 2, reveal intriguing correspondences with historical or political events. For instance, between '60s and '70s the study of physical properties of liquids was officially included in solid state physics, to form the basis of Condensed Matter, name adopted in that period to redirected into one common field those physicists who were previously working on simple and complex matter [18].

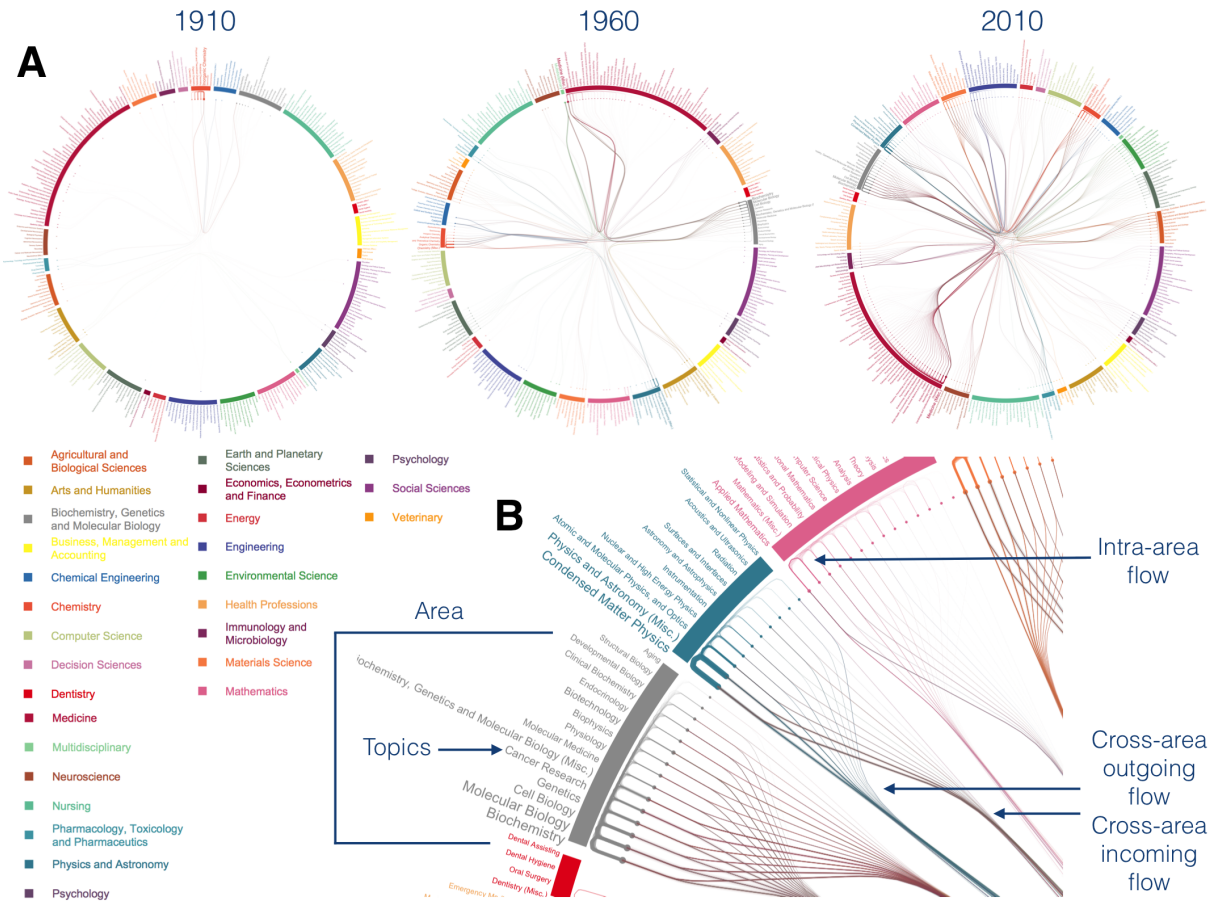

Figure 1. **Flow network of knowledge diaspora.** Points on the circle indicate topics (fine-grained knowledge representations) that are colored according to their SCImago area (coarse-grained knowledge representations), represented by thick sectors, whose color legend is reported. Two topics are connected if at least one author at time τ switched from one to another 5 years later. (A) Flow of authors moving his/her research activity from one topic to others across time. (B) How to read this visualization: switches between topics of the same area, namely "intra-area flows", are represented as 'U" shaped links close to sectors, to distinguish them from "cross-area flows". The outgoing flow is colored by the area of origin. The width of edges is proportional to the observed flow. See Appendix for more details about topics classification and this type of visualization.

Figure 2. **Most attractive topics in the knowledge diaspora.** The flow network of each temporal snapshot of 5 years is compared with the one immediately subsequent, and the relative changes in the volume of authors attracted by a topic (see Eq. (2)) are computed. For each temporal snapshot, we report the largest relative change observed in the volume. The relative increase is encoded in the size of circles.

Another interesting case is represented by Nanotechnology, with a significant activity change between 2000

and 2004, following the Nobel Prize in Chemistry won by Harry Kroto, Richard Smalley, and Robert Curl for the discovery of fullerenes. Fundamentals in many technological applications, fullerenes attracted a large number of researchers from –Statistics and Probability–, –Modeling and Simulation– and –Computer Science Applications–, when the new National Nanotechnology Initiative (`http://www.nano.gov/`) was officially proposed (1999) and the US President Bill Clinton declared a budget worth $500 million to support it (January 2000), thus justifying the diaspora from many other disciplines to –Nanoscience and Nanotechnology–. The case of –Agricultural and Biological Sciences (Misc.)–, exhibiting the largest value of $\delta_t(\tau)$ between 2010 and 2014, especially attracted our attention. A deeper analysis, revealed the presence of an increasing significant flow of researchers incoming from –Energy (Misc.)– who moved their publications towards in journals pertaining to agricultural and biological sciences, with research about genetically modified organisms, synthesis of biomolecules, biofuels, food systems and bioenergy.

After the fine-grained analysis at the level of topics, we focus on the analysis at the coarse-grained level of areas. For the analysis at the area level we need to define the intra-area flow as the volume of authors $V_a^{[\text{intra}]}(\tau)$ that keep publishing in the same area a over successive temporal snapshots. The overall cross-area incoming flow $V_a^{[\text{to}]}(\tau)$ is defined as the volume of authors who publish in area a at time τ coming from other areas. Finally, the overall cross-area outgoing flow $V_a^{[\text{from}]}(\tau)$ is defined as the volume of authors in area a that publish in other areas at time τ. These measures allow to investigate many aspects of the diaspora, characterizing the role played by different areas in the evolution of human knowledge. We introduce two local descriptors, namely the immigration and the emigration indices defined by

$$\iota_a(\tau) = \frac{V_a^{[\text{to}]}(\tau)}{V_a^{[\text{intra}]}(\tau) + V_a^{[\text{to}]}(\tau)} \tag{3}$$

$$\epsilon_a(\tau) = \frac{V_a^{[\text{from}]}(\tau)}{V_a^{[\text{intra}]}(\tau) + V_a^{[\text{from}]}(\tau)}, \tag{4}$$

respectively, characterizing the diaspora from a local perspective, i.e. in terms of relative variations with respect only to the existing population of authors working in the area a. These indices range from 0 – characterizing areas where the incoming (outgoing) flow of immigrating (emigrating) authors is negligible with respect to the existing authors population in the area – to 1 – indicating areas where the existing authors population is negligible with respect to the incoming (outgoing) flow of immigrating (emigrating) authors. However, these two local indices alone, do not allow to gain global insight about the diaspora from sources and to sinks of knowledge. For instance, such indices do not allow to understand if areas like –Physics and Astronomy–, –Mathematics– or –Computer Science–, producing academics whose modeling and abstraction skills make them suitable for challenging problems in other disciplines, act as global sources of the diaspora or not. In fact, it might happen that even if academics from these areas are commonly perceived to be very multidisciplinary, their flow with respect to the intra-area flow of authors could be rather small. To this aim we introduce two global descriptors, namely the sink and source indices defined by

$$\rho_a(\tau) = \frac{V_a^{[\text{to}]}(\tau)}{\sum_{a'} V_{a'}^{[\text{to}]}(\tau)} \tag{5}$$

$$\sigma_a(\tau) = \frac{V_a^{[\text{from}]}(\tau)}{\sum_{a'} V_{a'}^{[\text{from}]}(\tau)}, \tag{6}$$

respectively. As before, such indices range from 0 – indicating areas where the incoming (outgoing) flow of authors is negligible with respect to the overall incoming (outgoing) flow – to 1 – characterizing areas where the incoming (outgoing) flow of authors dominates the overall incoming (outgoing) flow.

In Fig. 3A–B is shown the evolution of the immigration and emigration across years for each area separately. Noticeably, most knowledge areas exhibit an evolution from an initial phase, where the incoming and outgoing flows of authors are negligible with respect to the existing authors population in the area, to the actual phase where these flows gain more and more importance. Nevertheless, some areas like Medicine, Physics and Astronomy, Chemistry or Mathematics are more secluded than others and partially preserve their isolation in both incoming and outgoing flow after one century. Conversely, a few areas like Nursing and Health Professions already exhibited a relevant outgoing flow almost a century ago. Of particular interest are those areas that were isolated a century ago but that, between '60s and '70s, have undergone a transition and started to both attract researchers from and provide researchers for other areas, such as Computer Science and Environmental Science. In the two decades between '50s and '70s many public and governmental research institutions invested on technological and theoretical investigation attracting, among others, mathematicians, physicists, philosophers and engineers. During the same years, the raise of Artificial Intelligence, required cross-disciplinary research at the edge of philosophy of mind, electrical engineering, neurophysiology, social intelligence and applied mathematics, to cite a few. In parallel, an inverse flow begun as well when a variety of disciplines started to take advantages of the new tools and methods provided by this area, like for example the emerging field of Digital Humanities. In the case of Environmental Science, the diaspora coincides with the revolution of the field in the '60s. In fact, the environmental movements born in that period to protest against chemical companies led to the creation of the U.S. Environmental Protection Agency and to the creation of many new environmental laws that required the development of specific environmental protocols of investigation, involving experts from a wide variety of disciplines. Fig. 3C

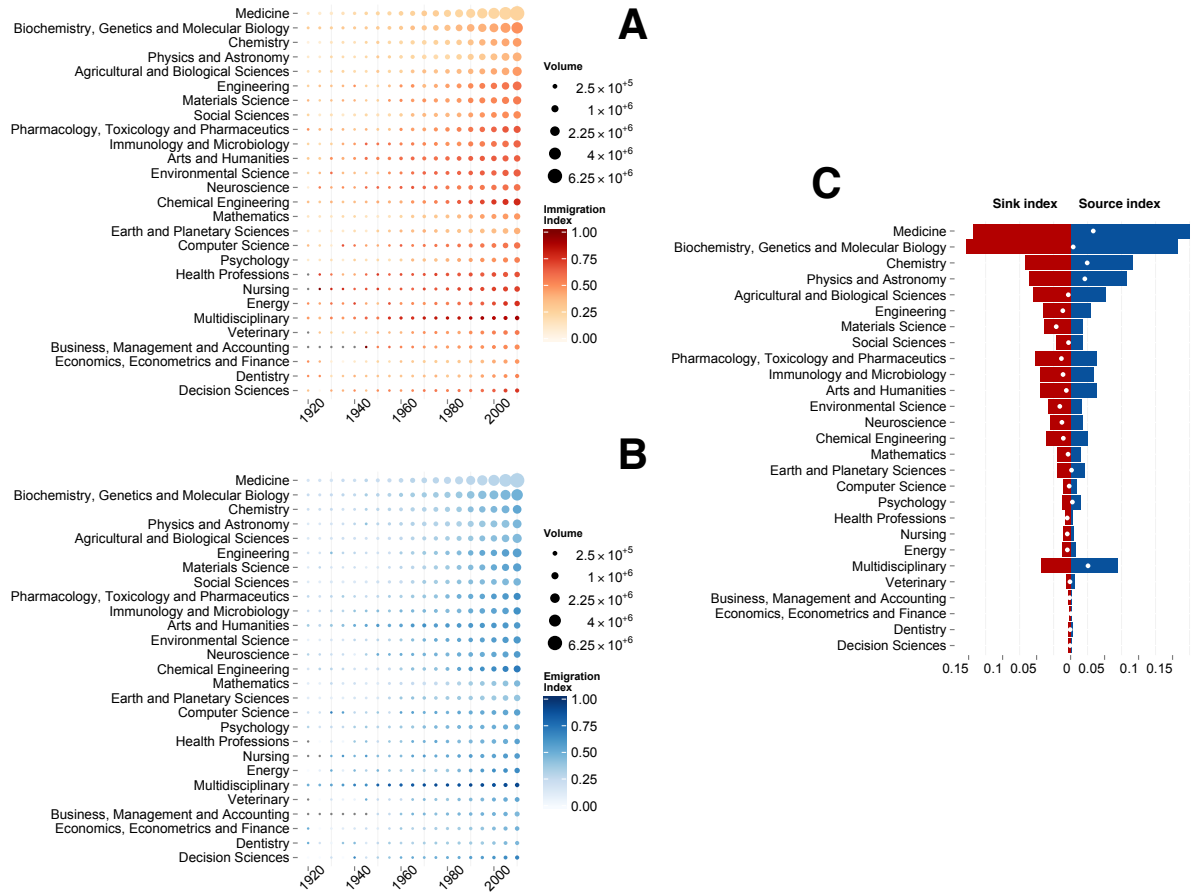

Figure 3. **Incoming and outgoing flows from and to knowledge areas across time.** Immigration (panel A) and emigration (panel B) index (see Eq. 2 and 3) of each area calculated for each temporal snapshot. Here, 0 indicates that the incoming (outgoing) flow of immigrating (emigrating) authors is negligible with respect to the existing authors population in the area, and 1 that the existing population is negligible with respect to the incoming (outgoing) flow of immigrating (emigrating) authors. Size of circles are proportional to the volume of authors in each area, and areas are ordered according to their overall volume over time. (C) Median sink (left, red boxes) and source (right, blue boxes) index (see Eq. 4 and 5) calculated for each area. Both range from 0 – indicating areas where the incoming (outgoing) flow of authors is negligible with respect to the overall incoming (outgoing) flow – to 1 – characterizing areas where the incoming (outgoing) flow of authors dominates the overall incoming (outgoing) flow.

shows the median over time of the source and sink indices for each area separately, which give instead a global perspective of incoming and outgoing flows. This allows to see that fields like Medicine and Physics, that seem isolated when analyzed locally, actually serve as sinks and sources of the knowledge diaspora. This means that, even though most research in these areas is carried out by authors who are already in the field, their contribution to the overall flow of knowledge is very relevant. In particular, both areas serve mostly as source of the diaspora, supplying other areas with researchers importing new methods and tools.

The knowledge diaspora obliged many researchers to work at the edge of different topics and different areas, driving an increasing trend towards higher transdisciplinary and multidisciplinary research, in agreement with very recent evidences [19]. Our data set allows us to quantify also the contribution of authors to different areas during the past 100 years. For each temporal snapshot of the network, we calculate the distribution of the number of different knowledge areas where an author has published in. The evolution of this distribution is shown in Fig. 4 where, as expected, we can observe how authors publish mainly in one area at the beginning of the past century while, over the years, a growing fraction of researchers has begun to produce publications in an increasing number of different areas.

V. DISCUSSION AND CONCLUSIONS

We have investigated the evolution of human knowledge across one century by using, as a proxy, the publication patterns of academics in different areas of research.

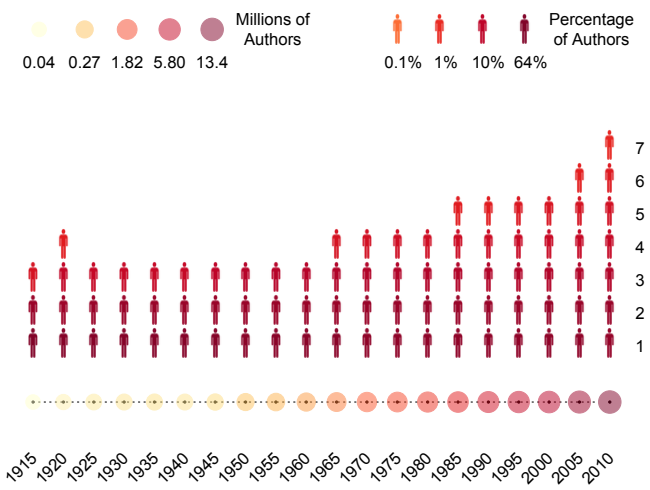

Figure 4. **Authors contribution to different areas.** Each column represents the 99%-quantile distribution of the number of different scientific areas that an author has published in during the corresponding temporal snapshot. Each icon represents, through its color, the density of authors having published in a given number of areas during a given temporal snapshot. The figure clearly shows that over time authors have increasingly started to publish in more and more scientific areas, i.e. they are becoming more and more multidisciplinary. The circles along the time axis represent the volume of authors that have published during the corresponding temporal snapshot, for reference.

For this purpose, we have used the Microsoft Academic Graph, the largest publicly available data set providing detailed information about academic publications in all areas of knowledge. Our multilayer network map allowed us to model the changes in research interests of academics across time, revealing what we called the "diaspora of the knowledge". In fact, we were able to identify disciplines acting as sources or sinks of academics' interest, quantifying their attractiveness across time and revealing fundamental periods in the raise of interest in areas of human knowledge. Noticeably, such periods might be related to crucial historical and political events. Our results show that, in the last century, a growing number of researchers published papers in an increasing number of disciplines. This clear trend illustrates, in a quantitative way, the perceived growth in the number of authors performing research crossing the boundaries of knowledge areas.

AUTHOR'S CONTRIBUTIONS

M.D.D and E.O. analyzed the data and performed the analysis. M.D.D., E.O. and A.A. designed the study and wrote the paper. All authors reviewed and approved the complete manuscript.

COMPETING INTERESTS

The authors declare no competing interests.

ACKNOWLEDGEMENTS

M.D.D. acknowledges financial support from the Spanish program Juan de la Cierva (IJCI-2014-20225). E.O. was supported by James S. McDonnell Foundation.. A.A. acknowledges financial support from ICREA Academia and James S. McDonnell Foundation and Spanish MINECO FIS2015-71582.

Appendix A: Building the diaspora network

Figure 5 illustrates how we define knowledge diaspora in terms of authors' movements across their research interests.

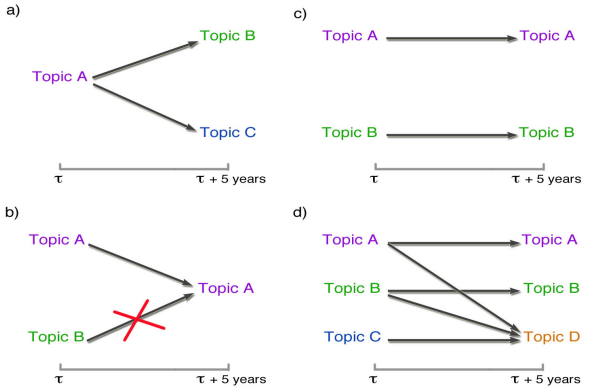

Figure 5. **Knowledge diaspora between areas.** (a) If an author publishes in different topics at time τ and at time $\tau + \Delta\tau$, we count one transition between all combinations of topics; (b) if an author publishes in topics A and B at time τ, and at time $\tau+\Delta\tau$ again in topic A but not in B anymore, then we consider just one self-transition from topic A to itself; (c) consistently, if an author publishes in topics A and B both at time τ and at time $\tau+\Delta\tau$, we only count two self-transitions; (d) more generally, if an author publishes in different topics at time τ, and one of them (C) disappears at time $\tau + \Delta\tau$, whereas another (D) appears, since we can not know from with topic at time τ there is a transition to topic D at time $\tau + \Delta\tau$, we therefore invoke the maximum entropy principle, suggesting that we have to count one transition from any topic at time τ to topic D.

Appendix B: Categorical edge-bundling visualization of networks

Visualizing in a clear and informative way the intricate web of transitions between different areas is a challenging problem. When the number of interested nodes, in

our case topics or areas, and their interconnections is sufficiently small, chord diagrams [20] are suitable candidates. However, if the number of interconnections is too large, chord diagrams might lose their high level of readability. We found a good alternative in edge-bundling visualization [21], although this approach requires hierarchical data and our network does not exhibit any natural hierarchy, that should instead obtained by applying external algorithms and it would be based on assumptions. Instead, what we wanted to exploit is the intrinsic categorization of authors and papers in areas and topics, while having full control on redirecting edges and place nodes according to our needing. Inspired by Circos visualization [22], we adopted a circular layout, i.e. embedding on a circle, where categories, in our case the areas of knowledge, are drawn as sectors with different colors. The position of sectors is chosen according to heuristics depending, among other factors, on the modular structure [23] of the network of layers. Nodes, in our case the topics, are placed on a circular layout, close to the sector encoding the area they belong to. Within each sector, nodes are ordered by the logarithm of their strength, to facilitate the identification of important topics and improve the visualization of connections. The size of nodes is rescaled to avoid nodes with radius below or above certain thresholds. The name of topics, i.e. node's label, is shown radially along the direction connecting the node to the center of the circle and both nodes and labels are colored according to the area they belong to, to facilitate readability. Edges are divided into three categories: "intra-area" (encoding connections among topics within the same area regardless of direction), "cross-area out-going" (encoding connections going from a topic to other topics in a different area) and "cross-area in-going" (encoding connections going to a topic from other topics in a different area). Intra-area edges are spline curves placed in the space between the sectors and the nodes, colored by using the color of the underlying area, allowing to gain insight about the diaspora within the same area. Cross-area edges are spline curves calculated by using five points, in addition to the positions of origin and destination nodes: 1) in front of the origin node, belonging to a "zero-level" circle with smaller radius than that where nodes are placed; 2) on a "first-level" circle with a smaller radius than the previous one, through a point whose position is on the right of the barycenter of the underlying sector and slightly displaced towards the sector; 3) on a "second-level" circle, through a point that is aligned with the barycenter of the two endpoints; 4) again on the first-level circle, through a point whose position is on the left of the barycenter of the underlying sector and slightly displaced towards the center of the circle, at variance with the third point; 5) one, on the zero-level circle, in front of the destination node. The displacement of a small angle to the right and to the left allows to separate the out-going and the in-going edges, respectively, as well as the small displacement along the radial direction facilitates the distinguishability of flow directionality, with out-going flow collected into a point closer to the sector and in-going flow collected into a point closer to the center of the circle. The color of each edge is calculated by interpolating the colors of the endpoints, while giving more weight to the color of the destination. The width of the edges is proportional to their weight, i.e. in our case the volume of authors between the endpoint topics and their transparency is regulated by the Euclidean distance between the connected nodes according to their position in the circular layout.

[1] K. W. Boyack and K. Börner, Journal of the American Society for Information Science and Technology **54**, 447 (2003).

[2] A. Ma, R. J. Mondragón, and V. Latora, Proceedings of the National Academy of Sciences p. In Press (2015).

[3] H. Etzkowitz and L. Leydesdorff, EASST Review **14**, 14 (1995).

[4] L. Leydesdorff and H. Etzkowitz, Science and public policy **25**, 195 (1998).

[5] H. Etzkowitz and L. Leydesdorff, Research policy **29**, 109 (2000).

[6] K. W. Boyack, R. Klavans, and K. Börner, Scientometrics **64**, 351 (2005).

[7] L. Leydesdorff and I. Rafols, Journal of the American Society for Information Science and Technology **60**, 348 (2009).

[8] P. Deville, D. Wang, R. Sinatra, C. Song, V. D. Blondel, and A.-L. Barabási, Scientific reports **4**, 4770 (2014).

[9] R. M. Shiffrin and K. Börner, Proceedings of the National Academy of Sciences **101**, 5183 (2004).

[10] K. Börner, J. T. Maru, and R. L. Goldstone, Proceedings of the National Academy of Sciences **101**, 5266 (2004).

[11] Q. Ke, E. Ferrara, F. Radicchi, and A. Flammini, Proceedings of the National Academy of Sciences p. 201424329 (2015).

[12] R. Sinatra, P. Deville, M. Szell, D. Wang, and A.-L. Barabási, Nature Physics **11**, 791 (2015).

[13] P. Holme and J. Saramäki, Phys. Rep. **519**, 97 (2012).

[14] M. Kivelä, A. Arenas, M. Barthelemy, J. P. Gleeson, Y. Moreno, and M. A. Porter, Journal of Complex Networks **2**, 203 (2014).

[15] M. De Domenico, A. Solè-Ribalta, E. Cozzo, M. Kivelä, Y. Moreno, M. A. Porter, S. Gòmez, and A. Arenas, Phys. Rev. X **3**, 041022 (2013).

[16] S. Boccaletti, G. Bianconi, R. Criado, C. Del Genio, J. Gómez-Gardeñes, M. Romance, I. Sendiña-Nadal, Z. Wang, and M. Zanin, Physics Reports **544**, 1 (2014).

[17] A. Sinha, Z. Shen, Y. Song, H. Ma, D. Eide, B.-j. P. Hsu, and K. Wang, in Proceedings of the 24th International Conference on World Wide We (International World Wide Web Conferences Steering Committee, 2015), pp. 243–246.

[18] J. D. Martin, Physics in Perspective **17**, 3 (2015).

[19] R. Van Noorden, Nature **525**, 306 (2015).

[20] G. J. Abel and N. Sander, Science **343**, 1520 (2014).

[21] D. Holten, Visualization and Computer Graphics, IEEE Transactions on **12**, 741 (2006).

[22] M. Krzywinski, J. Schein, I. Birol, J. Connors, R. Gascoyne, D. Horsman, S. J. Jones, and M. A. Marra, Genome research **19**, 1639 (2009).

[23] M. E. Newman, Nature Physics **8**, 25 (2012).